Inapproximability for VCG-Based Combinatorial Auctions

Dave Buchfuhrer* Shaddin Dughmi† Hu Fu‡ Robert Kleinberg§ Elchanan Mossel¶
Christos Papadimitriou‖ Michael Schapira** Yaron Singer†† Chris Umans‡‡

Abstract

The existence of incentive-compatible, computationally-efficient mechanisms for combinatorial auctions with good approximation ratios is the paradigmatic problem in algorithmic mechanism design. It is believed that, in many cases, good approximations for combinatorial auctions may be unattainable due to an inherent clash between truthfulness and computational efficiency. In this paper, we prove the first *computational-complexity* inapproximability results for incentive-compatible mechanisms for combinatorial auctions. Our results are tight, hold for the important class of VCG-based mechanisms, and are based on the complexity assumption that NP has no polynomial-size circuits. We show two different techniques to obtain such lower bounds: one for deterministic mechanisms that attains optimal dependence on the number of players and number of items, and one that also applies to a class of randomized mechanisms and attains optimal dependence on the number of players. Both techniques are based on novel VC dimension machinery.

1 Introduction

In a combinatorial auction, a set of items is sold to bidders with *private* preferences over *subsets* of the items, with the intent of maximizing the *social welfare* (*i.e.*, the sum of bidders' values for their allocated items). Manifesting the tension between bounded computational resources and strategic interaction between selfish participants, combinatorial auctions gained the status of being the paradigmatic problem in *algorithmic mechanism design* [31].

From a *computational perspective*, the general problem is NP-hard, and cannot be approximated within a constant factor [9]. From a *strategic perspective*, as agents' preferences are private, they may report false information in an attempt to manipulate the outcome. From a strictly computational perspective, extensive work in past years identified a rich class of instances that allow for positive computational results. While still NP-hard, the assumption that bidders' preferences are complement-free (the value for bundles does not exceed the sum of their components) allows for constant-factor approximations (see [9] for a survey). These approximations, however, assume agents reveal their *true* preferences. From a purely *strategic* perspective, the famous VCG mechanism can ensure bidders are incentivized to reveal their true preferences in this setting. This however, is under the assumption that one has unlimited computational resources, as the VCG mechanism requires the allocation problem to be solved *optimally* – an NP-hard task in our case.

At the heart of algorithmic mechanism design is the quest for auction protocols that are both *incentive-compatible* and *computationally efficient*, and guarantee decent approximation ratios. Sadly, to date, huge gaps exist between the state of the art approximation ratios obtained by unrestricted, and by truthful, algorithms.

*Computer Science Department at Caltech, Pasadena, CA, 91125 USA. Supported by NSF CCF-0346991, CCF-0830787 and BSF 2004329.

†Department of Computer Science, Stanford University. Email: shaddin@cs.stanford.edu. Most of this work was done while the author was visiting Cornell University. Supported by BSF grant 2006239, NSF grant CCF-0448664, and a Siebel Foundation Fellowship.

‡Department of Computer Science, Cornell University. Supported by NSF award CCF-0643934.

§Computer Science Department, Cornell University, Ithaca, NY 14853. Email: rdk at cs.cornell.edu. Supported by NSF awards CCF-0643934 and CCF-0910940, a Microsoft Research New Faculty Fellowship, and an Alfred P. Sloan Foundation Fellowship.

¶Statistics and Computer Science, U.C. Berkeley, and Mathematics and Computer Science Weizmann Institute. Supported by Sloan fellowship in Mathematics, NSF Career award DMS 0548249, DOD grant N0014-07-1-05-06, and by ISF. mossel@stat.berkeley.edu

‖Computer Science Division University of California at Berkeley, CA, 94720 USA. christos@cs.berkeley.edu

**Department of Computer Science, Yale University, CT, USA, and Computer Science Division, University of California at Berkeley, CA, USA. Supported by NSF grant 0331548. michael.schapira@yale.edu.

††Computer Science Division University of California at Berkeley, CA, 94720 USA. Supported by a Microsoft Research Fellowship. yaron@cs.berkeley.edu.

‡‡Computer Science Department at Caltech, Pasadena, CA, 91125 USA. Supported by NSF CCF-0346991, CCF-0830787 and BSF 2004329.

It is believed that this could be due to an inherent clash between the truthfulness and computational-efficiency requirements, that manifests itself in greatly degraded algorithm performance. For the first time, such tension between the two desiderata was recently shown to exist in [34] for a different mechanism design problem called *combinatorial public projects* [38]. However, in the context of combinatorial auctions, due to their unique combinatorial structure, *the algorithmic game theory community currently lacks the machinery to prove this* [36].

The celebrated class of *Vickrey-Clarke-Groves (VCG) mechanisms* [40, 11, 23] is the only known universal technique for the design of *deterministic* incentive-compatible mechanisms. (In certain interesting cases VCG mechanisms are the *only* truthful mechanisms [8, 20, 26, 35, 34].) While a naïve application of VCG is often computationally intractable, more clever uses of VCG are the key to the best known (deterministic) approximation ratios for combinatorial auctions [17, 24]. For these reasons, the exploration of the computational limitations of such mechanisms is an important research agenda (pursued in [16, 20, 26, 30, 34]). Recently, it was shown [34] that the computational complexity of VCG-based mechanisms is closely related to the notion of VC dimension. Using *existing* VC machinery, [34] was able to prove computational hardness results for combinatorial public projects. However, for combinatorial auctions, these techniques are no longer applicable. This is because, unlike combinatorial public projects, the space of outcomes in combinatorial auctions does not consist of subsets of the universe of items, but rather of *partitions* of this universe (between the bidders). This calls for the different VC machinery approaches for the handling of such problems.

1.1 Results In this paper, we show the first *computational complexity* lower bounds for VCG-based mechanisms, and truthful mechanisms in general, for combinatorial auctions (with the possible exception of a result in [26] for a related auction environment). First, we show this for deterministic maximal-in-range mechanisms for combinatorial auctions with budget-additive bidders. The class of budget-additive valuations, defined formally in Sec. 3, is strictly contained in the class of submodular valuations, which in turn is strictly contained in the class of complement-free valuations. Our inapproximability results depend on the computational assumption that SAT does not have polynomial-size circuits.

THEOREM 1.1. *Let $\mathcal{M}$ be a VCG-based mechanism in a combinatorial auctions with m items and n budget-additive bidders, where $n = n(m) \leq m^{\eta}$, for any positive constant $\eta < 1/2$. Then, $\mathcal{M}$ cannot approximate the social welfare within a factor better than $n/(1+\epsilon)$ unless $NP \subseteq P/poly$.*

This result is tight, as [17] show a VCG-based upper bound of $\sqrt{m}$, and a VCG-based upper bound of n is trivial.

Next, we extend our lower bound to a class of strictly more powerful randomized mechanisms. This class includes all universally-truthful VCG-based mechanisms, and more importantly a strictly more powerful class of truthful-in-expectation mechanisms — which we term *maximal-in-weighted-range (MIWR)*. It also includes every randomized mechanism that is a probability distributions over MIWR mechanisms; we call such mechanisms *randomized maximal-in-weighted-range*. Our result applies to any class of valuations satisfying natural closure properties — we term such valuation classes *regular* — and moreover rendering the algorithmic problem of two-bidder combinatorial auctions APX-hard. Such valuation classes include submodular, complement-free, superadditive, and coverage valuations, but not budget-additive valuations. The extension to randomized mechanisms comes at a cost, however: we show an optimal lower bound only in terms of the number of players. Nevertheless, as an easy corollary this result rules out a constant-factor approximation using randomized maximal-in-weighted-range mechanisms.

THEOREM 1.2. *Fix a regular valuation class $\mathcal{C}$ for which 2-player social welfare maximization is APX-hard. Fix a constant $n \geq 1$. For any constant $\epsilon > 0$, no polynomial-time randomized MIWR algorithm for n-player combinatorial auctions achieves an approximation ratio of $n - \epsilon$, unless $\mathsf{NP} \subseteq \mathsf{P/Poly}$.*

1.2 Techniques Informally, our method of lower bounding the approximability of deterministic VCG-based mechanisms via VC arguments is the following: We consider well-known auction environments for which *exact* optimization is NP-hard. We show that if a VCG-based mechanism *approximates closely* the optimal social welfare, then it is implicitly solving *optimally* a smaller, but still relatively large, optimization problem of the same nature — an NP-hard feat. We establish this by showing that the subset of outcomes (*partitions* of items) considered by the VCG-based mechanism is "large", and hence must "shatter" a relatively large subset of the items.

For the extension to randomized MIWR mechanisms, an additional idea is needed: rather than directly converting an r-approximate mechanism into an exact optimization algorithm for a smaller problem, we

instead convert an $(r + \delta)$-approximate MIWR algorithm into an $(r - \delta)$-approximate MIWR algorithm over a smaller set of items. When the problem is APX-hard, we then reach a contradiction by taking α to be the infimum of the approximation ratios achievable by polynomial-time MIWR mechanisms. Translating this idea into a rigorous proof requires a delicate induction over the number of players, as well as the careful handling of some complexity-theoretic difficulties.

For both results, dealing with partitions of items rather than subsets of items, as is traditional in VC-dimension machinery, poses a number of difficulties. For one, extending the Sauer-Shelah lemma to partitions requires carefully defining of what it means for a range to be "large": a lower-bound on the cardinality no longer suffices for a useful shattering result. Moreover, dealing with partitions requires that we handle the possibility of unallocated items. We overcome these difficulties by exhibiting counting arguments that show that, for any mechanism that obtains a non-trivial approximation ratio, there must be a reasonably large subset of items that are fully allocated in exponentially many different ways. This "shattering" of the allocation space allows us to establish our lower bounds.

1.3 Related Work Combinatorial auctions have been extensively studied in both the economics and the computer science literature [9, 12, 13]. It is known that if the preferences of the bidders are unrestricted then no constant approximation ratios are achievable (in polynomial time) [28, 33]. Hence, much research has been devoted to the exploration of restrictions on bidders' preferences that allow for good approximations, *e.g.*, for complement-free (*subadditive*), and *submodular*, preferences *constant* approximation ratios have been obtained [17, 19, 21, 22, 27, 41]. In contrast, the known *truthful* approximation algorithms for these classes have *non-constant* approximation ratios [14, 17, 18]. It is believed that this gap may be due to the computational burden imposed by the truthfulness requirement. However, to date, this belief remains unproven. In particular, no *computational complexity* lower bounds for truthful mechanisms for combinatorial auctions are known.

Vickrey-Clarke-Groves (VCG) mechanisms [40, 11, 23], named after their three inventors, are the fundamental technique in mechanism design for inducing truthful behaviour of strategic agents. Nisan and Ronen [30, 31] were the first to consider the computational issues associated with the VCG technique. In particular, [30] defines the notion of VCG-Based mechanisms. VCG-based mechanisms have proven to be useful in designing approximation algorithms for combinatorial auctions [17, 24]. In fact, the best known (deterministic) truthful approximation ratios for combinatorial auctions were obtained via VCG-based mechanisms [17, 24] (with the notable exception of an algorithm in [6] for the case that many duplicates of each item exist). Moreover, Lavi, Mu'alem and Nisan [26] have shown that in certain interesting cases VCG-based mechanisms are essentially the only truthful mechanisms (see also [20]).

Dobzinski and Nisan [16] tackled the problem of proving inapproximability results for VCG-based mechanisms by taking a *communication complexity* [42, 25] approach. Hence, in the settings considered in [16], it is assumed that each bidder has an *exponentially large string of preferences* (in the number of items). However, real-life considerations render problematic the assumption that bidders' preferences are exponential in size. Our intractability results deal with bidder preferences that are *succinctly described*, and therefore relate to *computational complexity*. Thus, our techniques enable us to prove lower bounds even for the important case in which bidders' preferences can be concisely represented.

The connection between the VC dimension and VCG-based mechanisms was observed in [34], where a general (i.e., not restricted to VCG-based mechanisms) inapproximability result was presented, albeit in the context of a different mechanism design problem, called *combinatorial public projects* (see also [38]). The analysis in [34] was carried out within the standard VC framework, and so it relied on *existing* machinery (namely, the Sauer-Shelah Lemma [37, 39] and its probabilistic version due to Ajtai [2]). To handle the unique technical challenges posed by combinatorial auctions (specifically, the fact that the universe of items is *partitioned* between the bidders) *new* machinery is required. Indeed, our technique can be interpreted as *an extension of the Sauer-Shelah Lemma to the case of partitions* in Sec. 4).

The VC framework has received much attention in past decades (see, *e.g.*, [3, 7, 29] and references therein), and many generalizations of the VC dimension have been proposed and studied. To the best of our knowledge, none of these generalizations captures the case of n-tuples of disjoint subsets of a universe considered in this paper. In addition, no connection was previously made between the the VC dimension and the approximability of combinatorial auctions.

1.4 Organization of the Paper In Sec. 2, we formally define the problem and develop the necessary technical background. In Sec. 3 we present our first result for maximal-in-range mechanisms as described above. In Sec. 4 we present our second result, pertaining to randomized maximal-in-weighted-range mechanisms.

We conclude and present open questions in Sec. 5.

2 Preliminaries

2.1 Combinatorial Auctions In a combinatorial auction there is a set $[m] = \{1, 2, \ldots, m\}$ of items, and a set $[n] = \{1, 2, \ldots, n\}$ of players. Each player i has a valuation function $v_i : 2^{[m]} \to \mathbb{R}^+$ that is normalized ($v_i(\emptyset) = 0$) and monotone ($v_i(A) \leq v_i(B)$ whenever $A \subseteq B$).

An allocation of items M to the players N is a function $S : M \to N \cup \{*\}$. Notice that we do not require all items to be allocated. If an allocation S allocates all items – i.e. S maps M into N – we say S is a *total allocation*. The allocation that allocates no items is called the *empty allocation*. For convenience, we use $S(j)$ to denote the player receiving item j, and we use S_i to denote the items allocated to player i. We use $\mathcal{X}(M, N)$ to denote the set of all alocations of M to N.

In combinatorial auctions, the feasible solutions are the allocations $\mathcal{X}([m], [n])$ of the items to the players. The *social welfare* of such an allocation S is defined as $\sum_i v_i(S_i)$. When the players have values $\{v_i\}_i$, we often use $v(S)$ as shorthand for the welfare of S. The goal in combinatorial auctions is to find an allocation that maximizes the social welfare.

2.2 Valuation Classes The hardness of designing truthful combinatorial auction mechanisms depends on the allowable player valuations. Recall that a *valuation* over M is a function $v : 2^M \to \mathbb{R}^+$. We let $\mathcal{V}$ denote the set of all valuations over all abstract finite sets M. A *valuation class* $\mathcal{C}$ is a subset of $\mathcal{V}$. Examples of valuation classes include submodular valuations, subadditive valuations, single-minded valuations, etc.

Our first hardness result pertains to deterministic mechanisms for a simple class: budget-additive valuations. This is despite the fact that the social welfare maximization problem admits an FPTAS when the number of bidders is constant [4]. Budget-additive valuations are defined as follows.

Definition 2.1. *We say a valuation $v : 2^M \to \mathbb{R}^+$ is* budget-additive *if there exists a constant $B \geq 0$ (the budget) such that $v(A) = \min(B, \sum_{i \in A} v(\{i\}))$.*

Our second result applies to any valuation class that induces an APX-hard welfare maximization problem over two bidders, and moreover satisfies some natural properties.

Definition 2.2. *We a say a valuation class $\mathcal{C}$ is* regular *if the following hold*

1. *Every valuation in $\mathcal{C}$ is monotone and normalized.*
2. *The canonical valuation on any singleton set is in $\mathcal{C}$. Namely, for any item a the valuation $v : 2^{\{a\}} \to \mathbb{R}^+$, defined as $v(\{a\}) = 1$ and $v(\emptyset) = 0$, is in $\mathcal{C}$.*
3. *Closed under scaling: Let $v : 2^M \to \mathbb{R}^+$ be in $\mathcal{C}$, and let $c \geq 0$. The valuation $v' : 2^M \to \mathbb{R}^+$, defined as $v'(A) = c \cdot v(A)$ for all $A \subseteq M$, is also in $\mathcal{C}$.*
4. *Closed under disjoint union: Let M_1 and M_2 be disjoint sets. Let the valuations $v_1 : 2^{M_1} \to \mathbb{R}^+$ and $v_2 : 2^{M_2} \to \mathbb{R}^+$ be in $\mathcal{C}$. Their disjoint union $v_3 = v_1 \oplus v_2 : 2^{M_1 \cup M_2} \to \mathbb{R}^+$, defined as $v_3(A) = v_1(A \cap M_1) + v_2(A \cap M_2)$ for all $A \subseteq M_1 \cup M_2$, is in $\mathcal{C}$.*
5. *Closed under relabeling: Let M_1, M_2 be sets with a bijection $f : M_2 \to M_1$. If $v_1 : 2^{M_1} \to \mathbb{R}^+$ is in $\mathcal{C}$, then the valuation $v_2 : 2^{M_2} \to \mathbb{R}^+$ defined by $v_2(S) = v_1(f(S))$ is also in $\mathcal{C}$.*

Note that all regular valuation classes support *zero-extension*. More formally, let $M \subseteq M'$, and let $v : 2^M \to \mathbb{R}^+$ be in $\mathcal{C}$. The extension of v to M', defined as $v'(A) = v(A \cap M)$ for all $A \subseteq M'$, is also in $\mathcal{C}$. In the context of combinatorial auctions, we use $\mathcal{C}_m$ to denote the subset of valuation class $\mathcal{C}$ that applies to items $[m]$.

Most well-studied valuation classes for which the underlying optimization problem is APX-hard are regular. This includes submodular, subadditive, coverage, and weighted-sum-of-matroid-rank valuations. However, in addition to budget-additive valuations, two interesting counter-examples come to mind: multi-unit (where items are indistinguishable), and single-minded valuations. Nevertheless, the underlying optimization problem is not APX hard for multi-unit valuations, and for single-minded valuations the computational hardness of approximation is $n^{1-\epsilon}$ even without the extra constraint of truthfulness (see [10]).

2.3 Truthfulness An n-bidder, m-item mechanism for combinatorial auctions with valuations in $\mathcal{C}$ is a pair (f, p) where $f : \mathcal{C}_m^n \to \mathcal{X}([m], [n])$ is an *allocation rule*, and $p = (p_1, \cdots, p_n)$ where $p_i : \mathcal{C}_m^n \to \mathbb{R}$ is a *payment scheme*. (f, p) might be either randomized or deterministic.

We say deterministic mechanism (f, p) is *truthful* if for all i, all v_i, v_i' and all v_{-i} we have that $v_i(f(v_i, v_{-i})_i) - p_i(v_i, v_{-i}) \geq v_i'(f(v_i', v_{-i})_i) - p(v_i', v_{-i})$. A randomized mechanism (f, p) is *universally truthful* if it is a probability distribution over truthful deterministic mechanisms. More generally, (f, p) is *truthful in expectation* if for all i, all v_i, v_i' and all v_{-i} we have that $E[v_i(f(v_i, v_{-i})_i) - p(v_i, v_{-i})] \geq E[v_i'(f(v_i', v_{-i})_i) - p_i(v_i', v_{-i})]$, where the expectation is taken over the internal random coins of the algorithm.

2.4 Algorithms and Approximation Fix a valuation class $\mathcal{C}$. An algorithm $\mathcal{A}$ for combinatorial auctions with $\mathcal{C}$ valuations takes as input the number of players n, the number of items m, and a player valuation profile $v_1, \ldots, v_n$ where $v_i \in \mathcal{C}_m$. $\mathcal{A}$ must then output an allocation of $[m]$ to $[n]$. For each n and m, $\mathcal{A}$ induces an allocation rule of m items to n bidders.

We say an algorithm $\mathcal{A}$ for n-player combinatorial auctions *achieves an α-approximation* if, for every input n, m and $v_1, \ldots, v_n$:

$$\alpha E[v(A(m, v_1, \ldots, v_n))] \geq \max_{S \in \mathcal{X}([m],[n])} v(S)$$

Moreover, we say $\mathcal{A}$ *achieves an α-approximation for m items* if the above holds whenever the number of items is fixed at m.

2.5 MIR, Randomized MIR, MIDR, and MIWR Maximal in range (MIR) algorithms were introduced in [32] as a paradigm for designing truthful approximation mechanisms for computationally hard problems. An algorithm $\mathcal{A}$ is maximal-in-range if it induces a maximal-in-range allocation rule when n and m are fixed.

DEFINITION 2.3. *An n-bidder, m-item allocation rule f is* maximal-in-range *(MIR) if there exists a set of allocations $\mathcal{R} \subseteq \mathcal{X}([m],[n])$, such that $\forall v_1, \ldots, v_n$ $f(v_1, ..., v_n) \in \arg\max_{S \in \mathcal{R}} \Sigma_i v_i(S_i)$.*

A generalization of maximal-in-range that uses randomization sometimes yields better algorithms. An algorithm $\mathcal{A}$ is *randomized maximal-in-range* if it induces a maximal-in-range allocation rule for every realization of its random coins. It is well known that a randomized MIR algorithm can be combined with the VCG payment scheme to yield a universally truthful mechanism.

Dobzinski and Dughmi defined a generalization of randomized maximal-in-range algorithms in [15], termed maximal-in-distributional-range (MIDR). Here, each element of the range is a distribution over allocations. The resulting mechanism outputs the distribution in the range that maximizes the expected welfare, and charges VCG payments.

DEFINITION 2.4. *f is* maximal-in-distributional-range *(MIDR) if there exists a set $\mathcal{D}$ of distributions over allocations such that for all $v_1, \ldots, v_n$, $f(v_1, ..., v_n)$ is a distribution $D \in \mathcal{D}$ that maximizes the expected welfare of a random sample from D.*

MIDR algorithms were used in [15] to obtain a polynomial-time, polynomial-communication, truthful-in-expectation FPTAS for multi-unit auctions, despite a lower bound of 2 on maximal-in-range algorithms that use polynomial communication. Moreover, they exhibited a variant of multi-unit auctions for which an MIDR FPTAS exists, yet no deterministic (or even universally truthful) polynomial time mechanism can attain an approximation ratio better than 2. Notably, the MIDR algorithms presented in [15] are of the following special form.

DEFINITION 2.5. *f is* maximal in weighted range *(MIWR) if f is MIDR, and moreover each distribution D in the range of f is a* weighted allocation*: There is a pure allocation $S \in \mathcal{X}([m],[n])$ such that D outputs S with some probability, and the empty allocation otherwise.*

We denote a weighted allocation that outputs S with probabiliby w by the pair (w, S). When there is room for confusion, we use the term *pure allocation* to refer to an unweighted allocation.

Our second result will apply to all distributions over MIWR mechanisms, a class of mechanisms we term *randomized maximal-in-weighted-range*. Randomized MIWR mechanisms include all randomized MIR mechanisms as a special case.

2.6 An MIR algorithm achieving a $\min(n, 2m^{1/2})$ approximation ratio Given valuation functions v_i for each bidder i, first form a bipartite graph with nodes on one side representing items and nodes on the other representing bidders. Form edges with weight $v_i(j)$ between the nodes representing bidder i and item j. Find a maximum weighted matching in this graph. Call the value of this matching $V_{matching}$. Now, consider $v_i([m])$, the value to player i of getting all the items. Let $V_{all} = \max_i v_i([m])$, and let i^* be the bidder that maximizes $v_i([m])$. If $V_{matching} \geq V_{all}$, assign items to bidders as in the maximum weighted matching. Otherwise, give every item to bidder i^*.

THEOREM 2.1. ([17], SLIGHTLY REPHRASED) *The above algorithm achieves a $\min(n, 2m^{1/2})$ approximation of the social welfare under subadditive valuations with free disposal.*

Proof. First, note that since there are n bidders, the maximum social welfare is at most n times the maximum welfare obtainable by any single bidder, V_{all}. So this algorithm is at most an n approximation. We now proceed to show that this algorithm is also at most a $2\sqrt{m}$ approximation.

Consider an assignment A which maximizes the social welfare. There are at most $\sqrt{m}$ bidders which get $\sqrt{m}$ or more of the items each. Call this set of bidders B_{high}, and call the others B_{low}.

If the bidders in B_{high} get more than half of the social welfare, V_{all} will be at least as great as the maximum value received by any bidder in B_{high}. Thus, V_{all} is at least $1/\sqrt{m}$ times the social welfare from bidders in B_{high}. Because the bidders in B_{high} get half the social welfare, the maximum social welfare is at most $2\sqrt{m}$ times V_{all} in this case.

In the other case, the bidders in B_{low} get at least half the social welfare. Consider the matching in the bidder-and-item graph in which every bidder in B_{low} receives the item maximizing $v_i(j)$ out of the items assigned to them in A. Since the valuations are subadditive and each bidder in B_{low} receives at most $\sqrt{m}$ items, the total value of B_{low} is at most $\sqrt{m}$ times the value of this matching. Since $V_{matching}$ is the maximal value over any matching, we see that the social welfare from B_{low} is at most $\sqrt{m}V_{matching}$. Thus, since B_{low} gets at least half the social welfare, the social welfare of A is at most $2\sqrt{m}$ times $V_{matching}$.

Since V_{all} is always an n approximation and one of $V_{all}, V_{matching}$ is a $2\sqrt{m}$ approximation of the social welfare, assigning items to achieve the max of these two welfares yields a $\min(n, 2\sqrt{m})$ approximation.

2.7 A Primer on Non-Uniform Computation

Non-uniform computation is a standard notion from complexity theory (see e.g. [5]). We say an algorithm is non-uniform if it takes in an extra parameter, often referred to as an *advice string*. However, the advice string is allowed to vary only with the size of the input (i.e. with m). Moreover, the length of the advice string can grow only polynomially in the size of the input. If a problem admits a non-uniform polynomial-time algorithm, this is equivalent to the existence of a family of polynomial-sized boolean circuits for the problem. When we say a non-uniform algorithm is polynomial-time MIR [MIWR], we mean that the algorithm runs in time polynomial in m, and maximizes over a [weighted] range, regardless of the advice string. When we say a non-uniform algorithm achieves an approximation ratio of α on m, we mean that there exists a choice of advice string for input length m such that the algorithm always outputs an α-approximate allocation.

2.8 Technical Assumptions For Second Result

For our second result, a note is in order on the representation of valuations. Our results hold in the computational model. Therefore, we may assume that valuation functions are *succinct*, in that they are given as part of the input, and can be evaluated in time polynomial in the length of their description. Naturally, this result applies to non-succint valuations with oracle access, when the resulting problem admits a suitable reduction from an APX-hard optimization problem.

Moreoever, due to the generality of this result, we need to make some technical assumptions. Namely, we restrict our attention to combinatorial auctions over a "well-behaved" family of instances. This restriction is without loss of generality for all well-studied classes of valuations for which the problem is APX-hard, such as coverage, submodular, etc. A family I of inputs to combinatorial auctions is *well-behaved* if there exists a polynomial $b(m)$ such that for each input $(k, m, v_1, \ldots, v_k) \in I$, the function v_i is represented as a bit-string of length $O(b(m))$, and moreover always evaluates to a rational number represented using $O(b(m))$ bits. While we believe this assumption may be removed, we justify it on two grounds. First, every well-studied variant of combinatorial auctions that is APX-hard is also APX-hard on a well-behaved family of instances, so this restriction is without loss for all such variants. Second, this assumption greatly simplifies our proof, since it allows us to describe the size of an instance by a single parameter, namely m.

3 Hardness for MIR Mechanisms

In this section, we prove the following theorem:

THEOREM 3.1. *Let $\mathcal{M}$ be a polynomial-time maximal-in-range mechanism for auctions with n budget-additive bidders and m items, with $n = n(m) \leq m^\eta$ for positive constant $\eta < 1/2$. If $\mathcal{M}$ approximates the social welfare with a ratio of $n/(1+\epsilon)$ for positive constant ϵ, then $NP \subseteq P/poly$.*

Theorem 3.1 is a direct consequence of Lemmas 3.2, 3.4 and 3.5 below. It also leads to the following theorem, which shows that it is not possible to find a polynomial-time maximal-in-range mechanism that achieves an approximation much better than the $\min(n, 2m^{1/2})$ in [17] unless NP has polynomial circuits.

THEOREM 3.2. *For any positive constant ϵ and $n = n(m) \leq poly(m)$, no polynomial-time maximal-in-range auction mechanism can approximate the social welfare with a ratio better than $\min(n, m^{1/2-\epsilon})$ by a constant factor unless $NP \subseteq P/poly$.*

Proof. This follows from Theorem 3.1 by simply noting that any mechanism $\mathcal{M}$ which performs well on $n = n(m) \leq m^{1/2-\epsilon}$ bidders will perform well on $n = n(m) \leq poly(m)$ bidders when all but $m^{1/2-\epsilon}$ of the bidders have valuation functions which are identically zero. Thus, by setting all but $m^{1/2-\epsilon}$ of the valuation functions to 0, and simulating $\mathcal{M}$, we are effectively simulating $\mathcal{M}$ on an auction with $n = m^{1/2-\epsilon}$, as assigning items to bidders with valuations functions

equal to zero has the same effect as not assigning them at all. Thus, setting $n' = \min(n, m^{1/2-\epsilon})$, we see by Theorem 3.1 that achieving an approximation ratio better than n' implies $NP \subseteq P/poly$.

We begin the proof of Theorem 3.1 by examining the structure of the range. Below we omit floors and ceilings when dealing with them would be routine.

3.1 The Counting Argument Let $\mathcal{M}$ be a maximal-in-range mechanism with range $R \subseteq ([n] \cup \{\star\})^m$. For a vector $x \in R$, $x_i = j$ means that item i is given to bidder j, while $x_i = \star$ indicates that no bidder is given item i. For $S \subseteq [m]$, we define R_S to be the subset of the range where all of the items in S are assigned to bidders,

$$R_S = \{x \in R \ : \ x_i \in [n] \text{ for all } i \in S\}.$$

When considering R_S we wish to focus on the bidders that the items in S are assigned to, so we define T_S to be the projection of R_S to the indices in S. So $T_S \subseteq [n]^{|S|}$.

In order to show that $\mathcal{M}$ can solve a hard problem, we will show that there is some T_S with sufficiently many elements so that subset sum can be embedded in the valuations of S by the various bidders in such a way that $\mathcal{M}$ will solve it. By focusing on a portion of the range such that there are no unassigned items within a fixed subset S, we can ignore the difficulties associated with unassigned items. This idea allows for the use of standard VC machinery. First, we show that there must be some exponentially large T_S. We begin with a helpful lemma.

LEMMA 3.1. *For any positive constant ϵ and any m, n for which the binomial coefficients below are positive,*

$$\frac{\binom{m}{\epsilon m/n}}{\binom{((1+2\epsilon)/n)m}{\epsilon m/n}} < \left(\frac{n}{1+\epsilon}\right)^{\epsilon m/n}.$$

Proof. First, note that

$$\frac{\binom{m}{\epsilon m/n}}{\binom{((1+2\epsilon)/n)m}{\epsilon m/n}} = \prod_{i=0}^{\epsilon m/n-1} \frac{m-i}{((1+2\epsilon)/n)m - i}.$$

Now,

$$\begin{aligned} \frac{m-i}{((1+2\epsilon)/n)m-i} &= \frac{m-i}{((1+2\epsilon)/n)m-i} \\ &< \frac{m}{(1+2\epsilon)m/n - \epsilon m/n} \\ &= \frac{n}{1+\epsilon} \end{aligned}$$

So multiplying the $\epsilon m/n$ terms together, we have

$$\begin{aligned} \frac{\binom{m}{\alpha m}}{\binom{((1+2\epsilon)/n)m}{\epsilon m/n}} &= \prod_{i=0}^{\epsilon m/n-1} \frac{m-i}{((1+2\epsilon)/n)m-i} \\ &< \prod_{i=0}^{\epsilon m/n-1} \frac{n}{1+\epsilon} \\ &= \left(\frac{n}{1+\epsilon}\right)^{\epsilon m/n}, \end{aligned}$$

which proves the lemma.

LEMMA 3.2. *Let $\mathcal{M}$ be a maximal-in-range mechanism for auctions with n bidders and m items that approximates the social welfare with a ratio of $n/(1+2\epsilon)$, for positive constant ϵ. Then there exists a set $S \subseteq [m]$ with $|S| = \epsilon m/n$ where T_S has size $|T_S| \geq (1+\epsilon)^{\epsilon m/n}$.*

Proof. To begin, we associate with each $x \in [n]^m$ a set of valuation functions. The valuation functions are such that

$$\begin{aligned} v_{i,j} &= \begin{cases} 1 & x_j = i \\ 0 & \text{otherwise} \end{cases} \\ b_i &= m. \end{aligned}$$

Let $x \in [n]^m$. Because $\mathcal{M}$ approximates the social welfare with a ratio of $(1+2\epsilon)/n$ and the maximum social welfare is m, there must be a member $r \in R$ of the range such that $r_i = x_i$ for at least $((1+2\epsilon)/n)m$ different indices i. Let S_x be the set of these indices,

$$S_x = \{i \ : \ r_i = x_i\}.$$

There are at least $\binom{|S_x|}{\epsilon m/n} \geq \binom{((1+2\epsilon)/n)m}{\epsilon m/n}$ subsets $S' \subseteq S_x$ of size $\epsilon m/n$. For each such set S', $T_{S'}$ contains the projection of x to S'. If $T_{S'}$ contains the projection of x to S', we say that x is *covered* by $T_{S'}$. If $t \in T_{S'}$ is the projection of x to S', we say that t *covers* x.

For a subset $S \subseteq [m]$, define $C(S)$ to be the number of vectors $x \in [n]^m$ which are covered by T_S. Since each $x \in [n]^m$ is covered by at least $\binom{((1+2\epsilon)/n)m}{\epsilon m/n}$ sets T_S with $|S| = \epsilon m/n$,

$$\sum_{S \subseteq [m], |S| = \epsilon m/n} C(S) \geq n^m \binom{((1+2\epsilon)/n)m}{\epsilon m/n}. \tag{3.1}$$

We now bound the sum $\sum_{S \subseteq [m], |S| = \epsilon m/n} C(S)$. Suppose by way of contradiction that for every subset $S \subseteq [m]$ of size $\epsilon m/n$, $|T_S| < (1+\epsilon)^{\epsilon m/n}$. Consider a subset $S \subset [m]$ such that $|S| = \epsilon m/n$. Each $t \in T_S$ covers $n^{m-\epsilon m/n}$ elements of $[n]^m$. So $C(S) < (1+\epsilon)^{\epsilon m/n} n^{m-\epsilon m/n}$, which

gives the bound

(3.2)
$$\sum_{S\subseteq[m],|S|=\epsilon m/n} C(S) < \binom{m}{\epsilon m/n}(1+\epsilon)^{\epsilon m/n} n^{m-\epsilon m/n}.$$

So by Equations 3.1 and 3.2, we have

$$\binom{m}{\epsilon m/n}(1+\epsilon)^{\epsilon m/n} n^{m-\epsilon m/n} > n^m \binom{((1+2\epsilon)/n)m}{\epsilon m/n},$$

which we simplify to

$$\frac{\binom{m}{\epsilon m/n}}{\binom{((1+2\epsilon)/n)m}{\epsilon m/n}}(1+\epsilon)^{\epsilon m/n} > n^{\epsilon m/n}. \tag{3.3}$$

By Lemma 3.1, we get

$$\frac{\binom{m}{\epsilon m/n}}{\binom{((1+2\epsilon)/n)m}{\epsilon m/n}}(1+\epsilon)^{\epsilon m/n} < \left(\frac{n}{1+\epsilon}\right)^{\epsilon m/n}(1+\epsilon)^{\epsilon m/n}$$

which is simply $n^{\epsilon m/n}$, contradicting (3.3). This proves that there exists some $S \subseteq [m]$ with $|S| = \epsilon m/n$ such that $|T_S| \geq (1+\epsilon)^{\epsilon m/n}$.

3.2 Using the VC Dimension At this point, we would like to use Sauer's lemma to show a large VC dimension. Unfortunately, it does not generalize well to auctions with three or more bidders because for $n > 2$ there exist sets of size $(n-1)^m > 2^m$ with n-ary VC dimension equal to 0. To get around this difficulty, we map T_S injectively from $[n]^{\epsilon m/n}$ into $[2]^{\epsilon m}$, and show that the image of this map has a large VC dimension. The large VC dimension then permits the embedding of an NP-hard problem (see Section 3.3). In order to show a lower-bound on the VC dimension, we use Sauer's Lemma:

LEMMA 3.3. (SAUER'S LEMMA) *Let S be a subset of $[2]^\ell$ with $|S| > \sum_{i=0}^{k-1}\binom{\ell}{i}$. The VC dimension of S is at least k.*

We will make use of the following corollary:

COROLLARY 3.1. *Let T be a subset of $[2]^\ell$. For any constant $\delta > 1/2$ and any $\epsilon > 0$, the following holds for all sufficiently large ℓ: if $|T| > (1+\epsilon)^{\epsilon\ell^\delta}$ then T has VC dimension at least $\ell^{1/2}$.*

Proof. Since for sufficiently large ℓ, $\ell^{1/2} < \ell/2$,

$$\begin{aligned}
\sum_{i=0}^{\ell^{1/2}-1}\binom{\ell}{i} &\leq \sum_{i=0}^{\ell^{1/2}-1}\binom{\ell}{\ell^{1/2}} \\
&\leq \ell^{1/2}\left(\frac{e\ell}{\ell^{1/2}}\right)^{\ell^{1/2}} \\
&= \ell^{1/2}\left(e\ell^{1/2}\right)^{\ell^{1/2}} \\
&= (1+\epsilon)^{1/2\log_{1+\epsilon}\ell+\ell^{1/2}\log_{1+\epsilon}(e\ell^{1/2})} \\
&= (1+\epsilon)^{\ell^{1/2}((1/2)\log_{1+\epsilon}\ell+\log_{1+\epsilon}e+o(1))} \\
&= (1+\epsilon)^{\ell^{1/2+o(1)}}
\end{aligned}$$

which is less than $|T| = (1+\epsilon)^{\epsilon\ell^\delta}$ for sufficiently large ℓ, since $\delta > 1/2$.

Let $\phi_i : n \to \{0,1\}$ be defined by

$$\phi_i(j) = \begin{cases} 1 & i = j \\ 0 & i \neq j \end{cases}.$$

For any vector v, take $\phi_i(v)$ to mean the application of ϕ_i to each component of v. Similarly, if ϕ_i is applied to a set of vectors, the result is that of applying ϕ_i to each vector in that set.

The next lemma is the main lemma in this section; it refers to the range R and the subsets T_S defined in Section 3.1.

LEMMA 3.4. *Let $\mathcal{M}$ be a maximal-in-range mechanism for auctions with n bidders and m items, with $n = n(m) \leq m^\eta$ for positive constant $\eta < 1/2$. For all sufficiently large m, if there exists a subset $S \subseteq [m]$ with $|S| = \epsilon m/n$ such that $|T_S| \geq (1+\epsilon)^{\epsilon m/n}$, then there exists a bidder i^* such that $\phi_{i^*}(R)$ has VC dimension at least $\sqrt{\epsilon}\cdot m^{1/2-\eta}$.*

Proof. Define vectors $e_j = (0,\ldots,0,1,0,\ldots,0)$, where the single 1 is in position j, and the number of coordinates of e_j is n. We define $f : [n]^{\epsilon m/n} \to [2]^{n\epsilon m/n} = [2]^{\epsilon m}$ by $f(x) = e_{x_1}e_{x_2}\cdots e_{x_{\epsilon m/n}}$. We write $f(T)$ for a subset T to mean the set $\{f(t) : t \in T\}$.

The function f is injective, so

$$|f(T_S)| = |T_S| \geq (1+\epsilon)^{\epsilon m/n}.$$

Note that $(1+\epsilon)^{\epsilon m/n} \geq (1+\epsilon)^{\epsilon m^{1-\eta}} \geq (1+\epsilon)^{\epsilon(\epsilon m)^{1-\eta}}$. We are assuming that m is sufficiently large, so we can apply Corollary 3.1 (with $\delta = 1-\eta > 1/2$ and $\ell = \epsilon m$) to conclude that $f(T_S)$ has VC dimension at least $(\epsilon m)^{1/2}$.

Let Q be a size $(\epsilon m)^{1/2}$ subset of $[\epsilon m]$ that is shattered by $f(T_S)$. Recall that each member of $f(T_S)$ is the concatenation of vectors of length n, where a 1 in

the ith position of the jth such vector corresponds to the ith bidder getting item j. In this way each element in Q corresponds to one of the n bidders. Partition Q into sets Q_i, where Q_i contains those coordinates that correspond to bidder i. There are n parts in the partition, so there is some $i^* \in [n]$ for which Q_{i^*} has size at least $(\epsilon m)^{1/2}/n$.

Since Q is shattered by $f(T_S)$, so is the subset Q_{i^*}. This means exactly that $\phi_{i^*}(T_S)$ has VC dimension at least $|Q_{i^*}| \geq (\epsilon m)^{1/2}/n$. Since the members of T_S are projections of members of R onto the coordinates in S, this implies that $\phi_{i^*}(R)$ also has VC dimension at least $(\epsilon m)^{1/2}/n \geq \sqrt{\epsilon} \cdot m^{1/2-\eta}$.

For a more intuitive understanding of Lemma 3.4, consider viewing all bidders other than i^* as a single meta-bidder. Lemma 3.4 states that there is a polynomially large set of items which are fully allocated in every possible way under this 2-bidder view.

3.3 Embedding Subset Sum We now show that if $\phi_{i^*}(R)$ has VC dimension at least m^γ for constant $\gamma > 0$, we can embed a subset sum instance into the auction in such a way that it is solved by $\mathcal{M}$. We use a reduction similar to one used in [27] to show that exactly maximizing the social welfare of these auctions is NP-hard.

LEMMA 3.5. *Let $\mathcal{M}$ be a polynomial-time maximal-in-range mechanism for auctions with n bidders and m items. Suppose there exists a constant $\gamma > 0$ such that for all sufficiently large m, there exists a bidder i^* such that $\phi_{i^*}(R)$ has VC dimension at least m^γ (where R is the range). Then $NP \subseteq P/poly$.*

Proof. We take as advice the set $L \subseteq [m]$ of size m^γ that is shattered by $\phi_{i^*}(R)$. For ease of exposition we re-order the items so that L is the set of the first m^γ items. Let $a_1, \ldots, a_{m^\gamma}$ be a subset sum instance with target sum K. For all bidders $i \neq i^*$, we set

$$\begin{aligned} v_{i,j} &= \begin{cases} a_j, & j \leq m^\gamma \\ 0, & j > m^\gamma \end{cases} \\ b_i &= \sum_j a_j \end{aligned}$$

and for bidder i^*, we set

$$\begin{aligned} v_{i^*,j} &= \begin{cases} 2a_j, & j \leq m^\gamma \\ 0, & j > m^\gamma \end{cases} \\ b_{i^*} &= 2K. \end{aligned}$$

If there is a subset V of $\{a_1, \ldots, a_{m^\gamma}\}$ summing to K, there is an assignment in R with social welfare of $\sum_j a_j + K$. This can be any assignment where bidder i^* gets the items in V, and the other items are distributed among the other bidders. R must contain such an assignment because $\phi_{i^*}(R)$ shatters L. Since $\mathcal{M}$ is maximal-in-range with range R, it will output an assignment with at least this social welfare.

If there is no subset of $\{a_1, \ldots, a_{m^\gamma}\}$ summing to K, $\mathcal{M}$ will assign bidder i^* a subset $V \subseteq M$ such that $\sum_{j \in V} a_j \neq K$. If $\sum_{j \in V} a_j < K$, the total value is at most

$$\begin{aligned} \sum_{j \notin V} a_j + \sum_{j \in V} 2a_j &= \sum_j a_j + \sum_{j \in V} a_j \\ &< \sum_j a_j + K. \end{aligned}$$

If $\sum_{j \in V} a_j > K$, bidder i^* gets value $2K$. So the total value is at most

$$\begin{aligned} \sum_{j \notin V} a_j + 2K &= \sum_j a_j - \sum_{j \in V} a_j + 2K \\ &< \sum_j a_j - K + 2K \\ &= \sum_j a_j + K. \end{aligned}$$

So every assignment has social welfare less than $\sum_j a_j + K$. So taking L as advice, we can solve a subset sum instance with k integers in polynomial time (in $m = k^{1/\gamma}$ and the size of the binary representations of the integers). Therefore, subset sum is in $P/poly$, so $NP \subseteq P/poly$.

3.4 Final Proof We can now prove Theorem 3.1. We have a polynomial-time maximal-in-range mechanism $\mathcal{M}$ for auctions with n bidders and m items, with $n = n(m) \leq m^\eta$ for positive constant $\eta < 1/2$. By Lemma 3.2, for each m there exists a subset $S \subseteq [m]$ of size $(\epsilon/2)m/n$ such that $|T_S| \geq (1+\epsilon/2)^{(\epsilon/2)m/n}$. By Lemma 3.4, this implies that for sufficiently large m, the range of $\mathcal{M}$ has VC dimension at least $\sqrt{\epsilon/2} \cdot m^{1/2-\eta}$. Since $\eta < 1/2$, we have $\sqrt{\epsilon/2} \cdot m^{1/2-\eta} \geq m^\gamma$ for some fixed positive constant γ and sufficiently large m. By Lemma 3.5, we thus have that $NP \subseteq P/poly$.

3.5 Super-polynomially many bidders In this section, we observe that our results can be extended to handle the case of n super-polynomial in m, at the expense of a stronger complexity assumption. For n larger than m, our technique shows a limit of $m^{1/2-\epsilon}$ on the approximation ratio of any mechanism which runs in time polynomial in m. However, by our definition an efficient mechanism need only run in time polynomial in n and

m, which is greater than $poly(m)$ for super-polynomial n. By strengthening the complexity assumption, we can still prove limits on the achievable social welfare.

For instance, if n is sub-exponential in m, we can begin by assuming that NP does not have sub-exponential size circuits. Then applying the same reduction leads to a circuit family of size $poly(n, m)$ (or sub-exponential in m), which solves subset sum instances of size m^γ for constant $\gamma > 0$, and this implies that NP has subexponential size circuits.

If n is sufficiently large as a function of m, it can even become possible to maximize the social welfare exactly in polynomial time.

THEOREM 3.3. *There exists a maximal-in-range mechanism $\mathcal{M}$ for auctions with n bidders and m items, which maximizes the social welfare and runs in polynomial time when $B_m \in O(poly(n))$, where B_m is the mth Bell number, the number of partitions of $[m]$ into any number of disjoint subsets with union $[m]$.*

Proof. If $B_m \in O(poly(n))$, it is possible to enumerate all of the partitions of $[m]$ into any number of disjoint subsets in polynomial time. For each such partition into sets $S_1, \ldots, S_k$, form a bipartite graph where one side has nodes representing the sets $S_1, \ldots, S_k$ and the other has nodes representing the bidders. The edge between bidder i and partition S_j has weight $v_i(S_j)$.

After finding a maximum weighted matching on each such bipartite graph, we can choose the maximum matching over all partitions. This matching represents the assignment which maximizes the social welfare. This can be easily seen because every assignment corresponds to a matching in the bipartite graph for some partition.

4 Hardness for Randomized MIWR Mechanisms

We find the proof in this section easier to follow when the approximation ratio is expressed by a number less than 1, and we will follow this practice henceforth. Thus, when we say that a randomized mechanism achieves an approximation ratio of $\beta < 1$, it means that for every input n, m and $v_1, \ldots, v_n$:

$$E[v(A(m, v_1, \ldots, v_n))] \geq \beta \cdot \left[\max_{S \in \mathcal{X}([m],[n])} v(S) \right].$$

In this section, we prove the following result.

THEOREM 4.1. *Fix a regular valuation class $\mathcal{C}$ for which 2-player social welfare maximization is APX-hard. Fix a constant $n \geq 1$. For any constant $\epsilon > 0$, no polynomial-time randomized MIWR algorithm for n-player combinatorial auctions achieves an approximation ratio of $1/n + \epsilon$, unless NP $\subseteq$ P/Poly.*

It is worth noting that this impossibility result applies to all universally-truthful randomized maximal-in-range algorithms. First,we prove the analogous result for MIWR mechanisms that take polynomial advice.

THEOREM 4.2. *Fix a regular valuation class $\mathcal{C}$ for which 2-player social welfare maximization is APX-hard. Fix a constant $n \geq 1$. For any constant $\epsilon > 0$, no non-uniform polynomial-time MIWR algorithm for n-player combinatorial auctions achieves an approximation ratio of $1/n + \epsilon$, unless NP $\subseteq$ P/Poly.*

We then complete the proof by showing that any randomized MIWR mechanism can be "de-randomized" to one that takes polynomial advice.

Our proof strategy for Theorem 4.2 is as follows. In Section 4.2 we define a "perfect valuation profile" on n players as a set of valuations where exactly one player is interested in each item. We then show that any range of allocations that gives a good approximation on a randomly drawn perfect valuation profile must "shatter" a constant fraction of the items, meaning that the range contains all allocations of that subset of the items to q of the players, where the value of q depends on the quality of the approximation.

In Section 4.3, we prove Theorem 4.2 by induction on the number of players n. Roughly speaking, we show that for any MIWR mechanism $\mathcal{A}$ for n players, the allocations with weight much larger than $1/n + \epsilon$ are useless. Namely, the inductive hypothesis implies that the allocations with weight sufficiently larger than $1/n+\epsilon$ cannot yield a good approximation to a randomly drawn perfect valuation; otherwise, one could use the resulting shattered set of items to design a strictly better MIWR mechanism for n' players for some $n' < n$. This allows us to conclude that all "useful" allocations have very similar weight to one another — the weights are close within $1-\eta$ for arbitrarily small η and a sufficiently large set of items. Now that the mechanism maximizes over a large set of allocations that are almost "pure", in the sense that the weights are almost identical, this yields a PTAS, contradicting the APX-hardness of the problem.

4.1 Some complexity theory. Before presenting the main body of the proof, we first develop some complexity theory tools to address a subtlety in the argument. Broadly speaking, our proof involves constructing a reduction that transforms every instance of a n-player mechanism design problem into an instance of one of n other problems $\mathcal{P}_1, \ldots, \mathcal{P}_n$, each of which individually is presumed to be computationally hard. The reduction has the property that input instances with a given number of items, m, are all transformed into inputs of

the same problem $\mathcal{P}_i$, but instances with a different number of items may map to a different one of the n problems. This raises difficulties because the complexity of $\mathcal{P}_1, \cdots, \mathcal{P}_n$ may be "wild": for each of them, there may be some input sizes (perhaps even infinitely many) that can be solved by a polynomial-sized Boolean circuit. In this section we develop the relevant complexity-theoretic tools to surmount this obstacle. We relegate the proofs of these results to Appendix B.

DEFINITION 4.1. *A set $S \subseteq \mathbb{N}$ is said to be* complexity-defying (CD) *if there exists a family of polynomial-sized Boolean circuits $\{\mathcal{C}_n\}_{n\in\mathbb{N}}$ such that for all $n \in S$, the circuit $\mathcal{C}_n$ correctly decides* 3SAT *on all instances of size n.*

A set $T \subseteq \mathbb{N}$ is said to be polynomially complexity-defying (PCD) *if there exists a complexity-defying set S and a polynomial function $p(n)$ such that T is contained in $\bigcup_{n\in S}[n, p(n)]$. Here $[a, b]$ denotes the set of all natural numbers x such that $a \leq x \leq b$. If a set $U \subseteq \mathbb{N}$ is not PCD, we say it is* non-PCD.

LEMMA 4.1. *A finite union of CD sets is CD, and a finite union of PCD sets is PCD.*

DEFINITION 4.2. *A decision problem or promise problem is said to have the* padding property *if for all $n < m$ there is a reduction that transforms instances of size n to instances of size m, running in time* $\mathrm{poly}(m)$ *and mapping "yes" instances to "yes" instances and "no" instances to "no" instances. Similarly, an optimization problem is said to have the padding property if for all $n < m$ there is a reduction that transforms instances of size n to instances of size m, running in time* $\mathrm{poly}(m)$ *and preserving the optimum value of the objective function.*

LEMMA 4.2. *Suppose that $\mathcal{L}$ is a decision problem or promise problem that has the padding property and is NP-hard under polynomial-time many-one reductions. Let T be any subset of $\mathbb{N}$. If there is a polynomial-sized circuit family that decides $\mathcal{L}$ correctly whenever the input size belongs to T, then T is PCD.*

LEMMA 4.3. *If $\mathbb{N}$ is PCD, then $\mathsf{NP} \subseteq \mathsf{P}/\mathrm{poly}$.*

4.2 Perfect Valuations and a strong shattering lemma We define a perfect valuation profile as one where each item is desired by exactly one player. Perfect valuation profiles will prove useful in our proof, due to the fact that no "small" range can well-approximate social welfare for a randomly-drawn perfect valuation profile.

DEFINITION 4.3. *Let N and M be a set of players and items, respectively. Let $v_i : 2^M \to \mathbb{R}^+$ be the valuation of player $i \in N$. We say the valuation profile $\{v_i\}_{i\in N}$ is a* perfect valuation profile *on N and M if there exists a total allocation S of M to N such that $v_i(j) = 1$ if $j \in S_i$, and $v_i(j) = 0$ otherwise. In this case, we say that $\{v_i\}_{i\in N}$ is the perfect valuation profile* generated *by S.*

To use perfect valuations in the proof of Theorem 4.1, they must belong to the class of valuations considered in the theorem. Indeed, it follows immediately from definition 2.2 that any regular class of valuations contains all perfect valuations.

The key property of perfect valuations is that, if a range $\mathcal{R}$ of allocations achieves a "good" approximation for many perfect valuations, then $\mathcal{R}$ must include all allocations of a constant fraction of the items to some set of q players. The value of q depends on the approximation guaranteed by $\mathcal{R}$, with a better approximation yielding a larger q. This is formalized by the following Lemma.

LEMMA 4.4. *Let U and V be finite sets with $|U| = m$, $|V| = n$, and R a set of functions from U to $V \cup \{*\}$. Suppose that for a random $f : U \to V$, with probability at least γ there is a $g \in R$ such that $g(x)$ differs from $f(x)$ on at most $\left(1 - \frac{q-1}{n} - \epsilon\right) m$ elements $x \in U$. Then there is a subset $S \subseteq U$ of cardinality at least δm (where $\delta > 0$ may depend on γ, ϵ, q, n) and a subset $T \subseteq V$ of cardinality q, such that every function from S to T occurs as the restriction of some $g \in R$.*

We note that this shattering lemma is more general than the Sauer-Shelah lemma, which can only be applied directly to two players. The notion of "shattering" among multiple players is key to the proof in Section 4.3.

If we interpret U as the set of items, V the set of players, f a perfect valuation profile, and R a range of allocations, then the implication of the lemma to our problem is immediate. The proof of the lemma is nontrivial, and we relegate it to Appendix A.

4.3 Hardness for Non-Uniform MIWR Mechanisms In this section we prove Theorem 4.2. We fix the valuation class $\mathcal{C}$ as in the statement of the theorem. Moreover, we fix $\eta > 0$ such that the 2-player social welfare maximization problem is NP-hard to approximate within a factor of $1 - \eta$. The proof proceeds by induction on n, the number of players. We need the following strong inductive hypothesis:

IH(n): *For any constant $\alpha > 1/n$ and set $T \subseteq \mathbb{N}$, if a non-uniform polynomial-time MIWR algorithm for the n-player problem achieves an α-approximation for m items whenever $m \in T$, then T is PCD.*

In other words, the set of input lengths for which any particular such algorithm may achieve an α-approximation is PCD. (See Section 4.1 for the definition of PCD.) If we can establish **IH(n)** for all $n \geq 1$, then Theorem 4.2 follows, since $\mathbb{N}$ is not PCD. The base case of $n = 1$ is trivial. We now fix n, and assume **IH(q)** for all $q < n$.

Assume for a contradiction that **IH(n)** is violated for some α. Let $\alpha > 1/n$ be the supremum over all values of α violating it. Note that **IH(n − 1)** implies that $\alpha \in \left(\frac{1}{n}, \frac{1}{n-1}\right]$. To simplify the exposition, we assume the supremum is attained, and fix the algorithm $\mathcal{A}$ (and corresponding family of polynomial advice strings) achieving an α-approximation for all $m \in \mathbb{F}$ where $\mathbb{F}$ is not PCD. Our arguments can all be easily modified to hold when the supremum is not attained, by instantiating $\mathcal{A}$ to achieve $(\alpha - \zeta)$ instead, where $\zeta > 0$ is as small as needed for the forthcoming proof. The proof then proceeds as follows. We assign every $m \in \mathbb{F}$ to one or more subsets $T_q\,(2 \leq q \leq n+1)$ in a way that is to be explained. We will prove that T_q is PCD for all q, and hence by Lemma 4.1 their union $\mathbb{F}$ is PCD. This, however, contradicts our assumption that $\mathbb{F}$ is not PCD and completes the proof.

To prove that T_q is PCD, we distinguish three cases depending on the value of q:

1. If $2 \leq q \leq n-1$, then we prove that $m \in T_q$ implies that there is a non-uniform polynomial-time MIWR mechanism for the q-player problem that achieves an approximation ratio strictly better than $1/q$ when the number of items is $\lceil \sigma m \rceil$, for some constant $\sigma > 0$. By our induction hypothesis, the set of all such $\lceil \sigma m \rceil$ must be a PCD set. By definition of PCD sets, this entails that T_q is PCD.

2. If $q = n$, then we prove that for the n-player problem, whenever $m \in T_q$, there is an efficient nonuniform MIWR mechanism that achieves approximation ratio strictly better than α for $\lceil \sigma m \rceil$ items. Once again this implies that the set of all such $\lceil \sigma m \rceil$ is a PCD set, by our hypothesis on α, which in turn implies that T_n is PCD.

3. If $q = n+1$ then we prove that when $m \in T_q$ there is a non-uniform polynomial-time algorithm achieving approximation ratio $1-\eta$ for the two-player m-item social welfare maximization problem. (Recall that η was chosen so that this problem is NP-hard to approximate within $1-\eta$.) By Lemma 4.2, this implies T_q is PCD.

We now implement this plan of proof in more detail.

Defining the partition of the range. Recall that an MIWR mechanism fixes a range of weighted allocations for each m. Let $\mathcal{D}^m$ denote the range of $\mathcal{A}$ when the number of items is m. Let $\mathcal{R}^m = \{S \in \mathcal{X}([m],[n]) : (w,S) \in \mathcal{D}^m \text{ for some } w\}$ be the corresponding set of pure allocations. For each allocation $S \in \mathcal{R}^m$, we use $w(S)$ to denote the weight of S in $\mathcal{D}^m$. (We assume without loss of generality that there is a unique choice of $w(S)$, since allocations with greater weight are always preferred.) When $m \in \mathbb{F}$, we may assume without loss of generality that $w(S) \geq \alpha$ for every $S \in \mathcal{R}^m$, since $\mathcal{A}$ achieves an α-approximation for m. We fix $\epsilon > 0$ such that $\alpha > 1/n + \epsilon$, and $\xi > 0$ such that $1/n + \epsilon/2 = (1-\xi)^{-1} \cdot (1/n)$, and let $\delta = \epsilon\eta/5n$. We partition $\mathcal{R}^m$ into *weight classes* as follows:

- $\mathcal{R}^m_q = \{S \in \mathcal{R}^m : \frac{1}{(1-\xi^2)q} \leq w(S) < \frac{1}{(1-\xi^2)(q-1)}\}$, for $2 \leq q \leq n-1$.
- $\mathcal{R}^m_n = \{S \in \mathcal{R}^m : \frac{\alpha}{1-\delta} \leq w(S) < \frac{1}{(1-\xi^2)(n-1)}\}$
- $\mathcal{R}^m_{n+1} = \{S \in \mathcal{R}^m : \alpha \leq w(S) < \frac{\alpha}{1-\delta}\}$

We partition $\mathcal{D}^m$ similarly: $\mathcal{D}^m_q = \{(w(S), S) : S \in \mathcal{R}_q\}$ for $2 \leq q \leq n+1$.

We are now ready to define T_q's. Roughly speaking, $m \in T_q$ if the output of $\mathcal{A}$, when applied to a random perfect valuation profile, has probability at least $1/n$ of being in $\mathcal{D}^m_q$.

Consider first the set $\mathcal{V}^m$ of perfect valuation profiles on $[n]$ and $[m'] = \{1, \ldots, m/2\}$, extended to $[m]$ by zero-extension. Given $v \in \mathcal{V}^m$ and $2 \leq q \leq n$, let us say that $v \in \mathcal{V}^m_q$ if the set $\mathcal{R}^m_q$ contains an allocation S that achieves at least a $(1+\xi)(q-1)/n$ approximation to the social welfare maximizer, and we say that $v \in \mathcal{V}^m_{n+1}$ if v does not belong to $\mathcal{V}^m_q$ for any $q < n+1$. Notice, first, that $\mathcal{V}^m_{n+1} \cap \mathcal{V}^m_q = \emptyset$ for every $2 \leq q \leq n$; second, that if $v \notin \mathcal{V}^m_q$, $(2 \leq q \leq n+1)$, then the best approximation ratio achievable using an allocation in $\mathcal{D}^m_q$ is at most

(4.4)
$$\frac{(1+\xi)(q-1)}{n} \cdot \frac{1}{(1-\xi^2)(q-1)} = \frac{1}{(1-\xi)n} = \frac{1}{n} + \frac{\epsilon}{2} < \alpha.$$

However, by our assumption that $\mathcal{A}$ achieves an α-approximation for all valuation profiles with $m \in \mathbb{F}$, the range $\mathcal{D}^m$ must contain an α-approximation to the social welfare maximizer. If $m \in \mathbb{F}$ and $v \in \mathcal{V}^m_{n+1}$, therefore, it follows that $\mathcal{D}^m_{n+1}$ must contain an α-approximation to the social welfare maximizer.

By the pigeonhole principle, at least one q satisfies

$$|\mathcal{V}^m_q| \geq \frac{1}{n} \cdot n^{m'}. \tag{4.5}$$

Finally, we define T_q to be the set of all $m \in \mathbb{F}$ such that (4.5) holds. By the preceding discussion, we have

$\mathbb{F} = \cup_{q=2}^{n+1} T_q$. We now proceed to prove that T_q is a PCD set for all q, completing the proof.

CLAIM 4.1. *T_q is a PCD set for all $q \in \{2, \ldots, n+1\}$.*

Cases 1 and 2: Large weight classes ($q \leq n$). To each allocation S of m items to n players, we may associate a function $f_S : [m'] \to [n] \cup \{*\}$, that maps each item $x \in [m']$ to the player who receives that item in S, or $*$ if the item is unallocated. Similarly, to each perfect valuation profile v on $[n]$ and $[m']$ we may associate a function $f_v : [m'] \to [n]$ that maps each item to the unique player who assigns a nonzero valuation to that item. Note that S achieves a c-approximation to the social-welfare-maximizing allocation for v if and only if the functions f_S and f_v differ on $(1-c)m'$ or fewer elements of $[m']$.

Assume now that $q \leq n$. If $m \in T_q$ then at least $1/n$ fraction of all perfect valuation profiles in $\mathcal{V}_q^m$ have an allocation $S \in \mathcal{R}_q^m$ that achieves a $(1+\xi)(q-1)/n$-approximation to the maximum social welfare. Thus, for at least $1/n$ fraction of all perfect valuation profiles $v \in \mathcal{V}_q^m$, there is some $S \in \mathcal{R}_q^m$ such that f_S and f_v differ on $\left(1 - \frac{q-1}{n} - \frac{(q-1)\xi}{n}\right) m'$ or fewer elements of $[m']$. Applying Lemma 4.4, there is a set W of at least $\lceil \sigma m \rceil$ elements of $[m']$, and a set N' of q players in $[n]$, such that all allocations of W to N' occur as restrictions of allocations in $\mathcal{R}_q^m$. We refer to W as a "shattered" subset of $[m']$.

When $q < n$ (Case 1 of our argument) we may now construct, via a non-uniform polynomial-time reduction, an MIWR allocation rule for the q-player problem that achieves a $[(1-\xi^2)q]^{-1}$ approximation for $\lceil \sigma m \rceil$ items when $m \in T_q$. Using W and N' – as defined above – as advice, embed the instance into an input for $\mathcal{A}$ by using players N' and items W in the obvious way: give player in $[n] \setminus N'$ an all-zero valuation. Moreover, extend the valuation of a player $i \in N'$ to the entire set of items $[m]$, assigning zero marginal value to every element of $[m] \setminus W$. Now, run $\mathcal{A}$ on the embedded instance. Notice that, since W is shattered by $\mathcal{R}_q^m$ with respect to N', every possible allocation of W to N' appears as the restriction of some allocation in R_q^m, and is therefore in the range of $\mathcal{A}$ with weight at least $[(1-\xi^2)q]^{-1}$. Thus, $\mathcal{A}$ must output a weighted allocation with expected welfare at least $[(1-\xi^2)q]^{-1}$ of the optimal. The result is a non-uniform poly-time MIWR mechanism for q players with approximation ratio bounded away from $1/q$ for all integers $\hat{m} = \lceil \sigma m \rceil$ such that $m \in T_q$. By our induction hypothesis **IH(q)**, this implies that the set of all such $\hat{m}$ is PCD. The fact that T_q itself is a PCD set now follows as an easy application of the definition of PCD.

When $q = n$ (Case 2 of our argument) using the same embedding yields an algorithm for n players that achieves an $\alpha/(1-\delta)$ approximation for all $\hat{m} = \lceil \sigma m \rceil$ such that $m \in T_q$. By our definition of α, this implies that the set of all such $\hat{m}$ is a PCD set, which again implies that T_q is a PCD set.

Case 3: The smallest weight class. The remaining case is $q = n+1$. When $m \in T_{n+1}$, by our definition of $\mathcal{V}_{n+1}^m$, at least $1/n$ fraction of all (extended) perfect valuation profiles $v \in \mathcal{V}^m$ have a weighted allocation $(w(S), S) \in \mathcal{D}_{n+1}^m$ that is an α-approximation to the social welfare maximizing allocation for v. Since $\alpha \leq w(S) < \alpha/(1-\delta)$, the pure allocation S must be a $(1-\delta)$-approximation to the social welfare maximizer. On the other hand, our assumption is that maximizing social welfare is APX-hard, even for two players; specifically, recall that $\eta > 0$ was chosen such that it is NP-hard to approximate the maximum social welfare within a factor of $1-\eta$. We complete the proof by exhibiting a randomized, non-uniform polynomial time algorithm that achieves a $(1-\eta)$-approximation for the n-player problem with $m/2$ items, for all $m \in T_{n+1}$. Notice that the de-randomization argument of Adleman [1] for proving $\mathsf{BPP} \subseteq \mathsf{P/Poly}$ can be used to de-randomize this to a non-uniform deterministic $(1-\eta)$-approximation for the n-player problem with $m/2$ items, for all $m \in T_{n+1}$. The reader unfamiliar with Adleman's argument may refer to Section 4.4, where we use the argument to establish Theorem 4.1.

We will now use $\mathcal{A}$ to get a $(1-\eta)$-approximate solution for an instance with n players and $m' = m/2$ items for all $m \in T_{n+1}$. We embed the instance on n players and m' items into $\mathcal{A}$ in the following way. Let M_1 be $[m] \setminus [m']$ and $v_i : 2^{M_1} \to \mathbb{R}$ the valuation function of player i. We assume without loss of generality that $\max_i v_i(M_1) = 1$. Next, we modify each player's valuation function by "mixing in" a perfect valuation profile on the remaining set of items $M_2 = [m']$. We draw a perfect valuation profile $(v_1', \ldots, v_n')$ on N and M_2 uniformly at random. Now, we "mix" the original valuations v with v', in proportions 1 and $\gamma = \frac{4n}{\epsilon m'}$, to yield the following *hybrid valuation profile* $v^* : 2^M \to \mathbb{R}^+$.

$$v_i^* = v_i \oplus \gamma v_i'$$

We abuse notation and use v_i $[v_i']$ to refer also to the zero-extension of v_i $[v_i']$ to M. Let OPT, OPT', and OPT^* be the optimal social welfare for the valuation profiles $\{v_i\}$, $\{v'\}$ and $\{v^*\}$, respectively. Then $1 \leq OPT \leq n$, and $OPT' = m'$, by construction. Since v and v' are defined on two disjoint sets of items, it is easy to see that $OPT^* = OPT + \gamma OPT'$. The scalar γ was carefully chosen so that the following facts hold:

1. The random valuation profile v' accounts for a majority share of v^* in any optimal solution. Specif-

ically, $\gamma OPT' \geq \frac{4}{\epsilon} OPT$. This implies that an allocation that gives a good approximation to OPT^* gives a similar approximation to OPT'. To be more precise, it can be shown by a simple calculation that:

CLAIM 4.2. *For any $S \in \mathcal{X}$ and any $\beta \geq 0$, if $v^*(S) \geq \beta OPT^*$ then $v'(S) \geq (\beta - \epsilon/2)OPT'$.*

2. The original valuation profile v accounts for a constant-factor share of v^* in any optimal solution. Specifically $OPT \geq \frac{\epsilon}{4n}(\gamma OPT')$. This implies that an allocation that gives $(1-\delta)$-approximation to OPT^* gives a $(1 - O(\delta))$-approximation to OPT. To be more precise, it can be shown by a simple calculation that:

CLAIM 4.3. *For any $S \in \mathcal{X}$, if $v^*(S) \geq (1-\delta)OPT^*$ then $v(S) \geq (1-\frac{5n}{\epsilon}\delta)OPT = (1-\eta)OPT$. (Recall that $\delta = \epsilon\eta/5n$.)*

We are now ready to show that running $\mathcal{A}$ on the valuations v^* will yield, with constant probability, an allocation that is a $(1-\eta)$-approximation to the optimal welfare for the original valuations v, when $m \in T_{n+1}$. Let $(w(S), S)$ be the weighted allocation output by $\mathcal{A}$; note that S is a random variable over draws of v'. Since $\mathcal{A}$ is an α approximation algorithm, the welfare $w(S)v^*(S)$ is at least $\alpha OPT^* \geq (1/n + \epsilon)OPT^*$ with probability 1. This implies that $v^*(S) \geq \left(\frac{1}{w(S)\cdot n} + \frac{\epsilon}{w(S)}\right) OPT^*$. By Claim 4.2, we see that $v'(S)$ is not too far behind: $v'(S) \geq \left(\frac{1}{w(S)\cdot n} + \frac{\epsilon}{w(S)} - \frac{\epsilon}{2}\right) OPT'$. Moreover, this gives:

$$w(S)v'(S) \geq \left(\frac{1}{n} + \frac{\epsilon}{2}\right) OPT' \tag{4.6}$$

Recall from equation (4.4) that if $v' \in \mathcal{V}^m_{n+1}$ then for $2 \leq q \leq n$, there is no $S \in \mathcal{R}^m_q$ that satisfies (4.6), and any such S satisfying (4.6) must belong to $\mathcal{D}^m_{n+1}$. Also, by our assumption that $m \in T_{n+1}$, the probability that $v' \in V^m_{n+1}$ is at least $1/n$.

We have thus established that running $\mathcal{A}$ on the random input v^* yields, with probability at least $1/n$, an outcome $(w(S), S)$ in $\mathcal{D}^m_{n+1}$. Using the fact that $w \leq \alpha/(1-\delta)$ and $w(S)v^*(S) \geq \alpha OPT$, we conclude that S is $(1-\delta)$-approximate for v^* also with probability $1/2$:

$$v^*(S) \geq (1-\delta)OPT^*$$

Invoking Claim 4.3, we conclude that $v(S) \geq (1-\eta)OPT$ with constant probability over draws of v'. Since $w(S)$ is at least $1/n$, S is output by $\mathcal{A}$ with constant probability. This completes the proof.

4.4 De-Randomizing Theorem 4.1 In this section, we complete the proof of Theorem 4.1. First, we make the observation that running a randomized MIWR algorithm multiple times independently and returning the best allocation output by any of the runs results in another randomized MIWR algorithm.

LEMMA 4.5. *Fix a randomized MIWR algorithm $\mathcal{A}$ and a positive integer r. Let $\mathcal{A}^r$ be the algorithm that runs r independent executions of $\mathcal{A}$ on its input, and of the r allocations returned, outputs the one with greatest welfare. $\mathcal{A}^r$ is also randomized MIWR.*

Proof. Condition on $\mathcal{D}_1, \ldots, \mathcal{D}_r$, the ranges of $\mathcal{A}$ on the r independent executions. $\mathcal{A}$ maximizes expected welfare over $\mathcal{D}_i$ on execution i. Therefore $\mathcal{A}^r$ maximizes over $\mathcal{D}_1 \cup \cdots \cup \mathcal{D}_r$.

Now, we derive Theorem 4.1 from Theorem 4.2, using a de-randomization argument similar to that of Adleman [1]. Assume for a contradiction that $\mathcal{A}$ is a randomized MIWR algorithm that runs in polynomial time and achieves an approximation ratio $1/n + \epsilon$ for each input m and $v_1, \ldots, v_n$. Let ℓ denote the number of bits in the input, and let $s(\ell)$ be a polynomial bounding the length of the random string drawn by $\mathcal{A}$. We will describe a polynomial-time with polynomial-advice MIWR algorithm that achieves an approximation ratio of $1/n + \epsilon/2$, which contradicts Theorem 4.2.

Let $r(\ell) = 2\ell/\epsilon^2$ and let $\mathcal{A}' = \mathcal{A}^{r(\ell)}$. By Lemma 4.5, $\mathcal{A}'$ is randomized MIWR, runs in polynomial time, and draws at most $s(\ell)r(\ell)$ random bits. Let X_i be the fraction of the optimal social welfare achieved by the allocation output on the i'th run of $\mathcal{A}$. The random variables $X_1, \ldots, X_{r(\ell)}$ are independent, $0 \leq X_i \leq 1$, and $E[X_i] \geq 1/n + \epsilon$. For each input of length ℓ, the probability that none of the $r(\ell)$ runs of $\mathcal{A}$ return an allocation with welfare better than $1/n + \epsilon/2$ of the optimal can be bounded from above using Hoeffding's inequality:

$$\begin{aligned} & Pr\left[\max_i X_i \leq \frac{1}{n} + \frac{\epsilon}{2}\right] \\ \leq\ & Pr\left[E\left(\sum_i X_i\right) - \sum_i X_i \frac{\epsilon r(\ell)}{2}\right] \\ \leq\ & e^{-\epsilon^2 r(\ell)/2} = e^{-\ell}. \end{aligned}$$

The number of different inputs of length ℓ is 2^ℓ. Thus, using the union bound and the above inequality, the probability that $\mathcal{A}$ outputs a $(1/n + \epsilon/2)$-approximate allocation on all inputs of length ℓ is non-zero. Therefore, for each ℓ there is choice of at most

$s(\ell)r(\ell)$ random bits such that $\mathcal{A}'$ achieves a $1/n + \epsilon/2$ approximation for all inputs. Using this as the advice string, this contradicts Theorem 4.2, completing the proof of Theorem 4.1.

5 Conclusions

We have shown that no polynomial-time maximal-in-range auction mechanism can approximate the social welfare to a ratio better than $\min(n, m^{1/2-\epsilon})$ by a constant factor. This essentially resolves the maximum social welfare achievable by efficient maximal-in-range auction mechanisms for any class of valuations including the valuation functions we considered, as a $\min(n, 2m^{1/2})$ ratio is achievable.

There is an asymmetry as to the strength of the n and $m^{1/2-\epsilon}$ bounds, however, as the n bound eliminates the possibility of a ratio of $n/(1+\epsilon)$ being achieved, but the $m^{1/2-\epsilon}$ bound leaves open the possibility of achieving an $m^{1/2-o(1)}$ approximation.

For super-polynomial n, we have demonstrated similar limits under stronger complexity assumptions, up to n being sub-exponential in m. We also showed that for sufficiently large n, a polynomial-time maximal-in-range auction mechanism exists.

Generalizing to randomized maximal-in-weighted-range mechanisms, we showed that it is impossible to achieve an approximation ratio better than n for any fixed n. In order to achieve these results, we developed new machinery for the study of the VC dimension of partitions. This new machinery allows for the application of a useful generalization of the standard VC dimension, and is therefore of independent interest.

While this largely resolves the performance of maximal-in-range and maximal-in-weighted-range mechanisms, it leaves open the performance of maximal-in-distributional-range mechanisms, as well as the larger question of how well truthful mechanisms perform.

References

[1] Leonard Adleman. Two theorems on random polynomial time. In *SFCS '78: Proceedings of the 19th Annual Symposium on Foundations of Computer Science*, pages 75–83, Washington, DC, USA, 1978. IEEE Computer Society.

[2] Miklós Ajtai. The shortest vector problem in L_2 is NP-hard for randomized reductions (extended abstract). In *STOC*, pages 10–19, 1998.

[3] Noga Alon, Shai Ben-David, Nicolò Cesa-Bianchi, and David Haussler. Scale-sensitive dimensions, uniform convergence, and learnability. *J. ACM*, 44(4):615–631, 1997.

[4] Nir Andelman and Yishay Mansour. Auctions with budget constraints. In Torben Hagerup and Jyrki Katajainen, editors, *SWAT*, volume 3111 of *Lecture Notes in Computer Science*, pages 26–38. Springer, 2004.

[5] Sanjeev Arora and Boaz Barak. *Computational Complexity: A Modern Approach*. Cambridge University Press, 1 edition, April 2009.

[6] Yair Bartal, Rica Gonen, and Noam Nisan. Incentive compatible multi unit combinatorial auctions. In *TARK '03: Proceedings of the 9th conference on Theoretical aspects of rationality and knowledge*, pages 72–87, New York, NY, USA, 2003. ACM.

[7] Peter L. Bartlett and Philip M. Long. Prediction, learning, uniform convergence, and scale-sensitive dimensions. *J. Comput. Syst. Sci.*, 56(2):174–190, 1998.

[8] Sushil Bikhchandani, Shurojit Chatterji, Ron Lavi, Ahuva Mu'alem, Noam Nisan, and Arunava Sen. Weak monotonicity characterizes deterministic dominant-strategy implementation. *Econometrica*, 74(4):1109–1132, 07 2006.

[9] Liad Blumrosen and Noam Nisan. Combinatorial auctions. In Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay Vazirani, editors, *Algorithmic Game Theory*, chapter 11. Cambridge University Press, 2007.

[10] Liad Blumrosen and Noam Nisan. Combinatorial auctions. In Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay Vazirani, editors, *Algorithmic Game Theory*, chapter 11. Cambridge University Press, 2007.

[11] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, pages 17–33, 1971.

[12] P. Cramton, Y. Shoham, and R. Steinberg (Editors). *Combinatorial Auctions*. MIT Press., 2006.

[13] Sven de Vries and Rakesh Vohra. Combinatorial auctions: A survey. *INFORMS jounral of computing*, 15(3):284–309, 2003.

[14] Shahar Dobzinski. Two randomized mechanisms for combinatorial auctions. In Moses Charikar, Klaus Jansen, Omer Reingold, and José D. P. Rolim, editors, *APPROX-RANDOM*, volume 4627 of *Lecture Notes in Computer Science*, pages 89–103. Springer, 2007.

[15] Shahar Dobzinski and Shaddin Dughmi. On the power of randomization in algorithmic mechanism design. In *FOCS '09: Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science*, Atlanta, GA, USA, 2009. To appear.

[16] Shahar Dobzinski and Noam Nisan. Limitations of vcg-based mechanisms. In *STOC '07: Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 338–344, New York, NY, USA, 2007. ACM.

[17] Shahar Dobzinski, Noam Nisan, and Michael Schapira. Approximation algorithms for combinatorial auctions with complement-free bidders. In *STOC '05: Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 610–618, New York, NY, USA, 2005. ACM.

[18] Shahar Dobzinski, Noam Nisan, and Michael Schapira. Truthful randomized mechanisms for combinatorial auctions. In *STOC '06: Proceedings of the thirty-*

eighth annual ACM symposium on Theory of computing, pages 644–652, New York, NY, USA, 2006. ACM.

[19] Shahar Dobzinski and Michael Schapira. An improved approximation algorithm for combinatorial auctions with submodular bidders. In *SODA '06: Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pages 1064–1073, New York, NY, USA, 2006. ACM.

[20] Shahar Dobzinski and Mukund Sundararajan. On characterizations of truthful mechanisms for combinatorial auctions and scheduling. In *EC '08: Proceedings of the 9th ACM conference on Electronic commerce*, pages 38–47, New York, NY, USA, 2008. ACM.

[21] Uriel Feige. On maximizing welfare when utility functions are subadditive. In *STOC '06: Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 41–50, New York, NY, USA, 2006. ACM.

[22] Uriel Feige and Jan Vondrak. Approximation algorithms for allocation problems: Improving the factor of 1 - 1/e. In *FOCS '06: Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science*, pages 667–676, Washington, DC, USA, 2006. IEEE Computer Society.

[23] T. Groves. Incentives in teams. *Econometrica*, pages 617–631, 1973.

[24] Ron Holzman, Noa Kfir-Dahav, Dov Monderer, and Moshe Tennenholtz. Bundling equilibrium in combinatrial auctions. *Games and Economic Behavior*, 47:104–123, 2004.

[25] Eyal Kushilevitz and Noam Nisan. *Communication Complexity*. Cambridge University Press, 1997.

[26] Ron Lavi, Ahuva Mu'alem, and Noam Nisan. Towards a characterization of truthful combinatorial auctions. In *FOCS '03: Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science*, page 574, Washington, DC, USA, 2003. IEEE Computer Society.

[27] Benny Lehmann, Daniel Lehmann, and Noam Nisan. Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behaviour. (A preliminary version appeared in EC'01)*, 55(2):270–296, 2006.

[28] Daniel Lehmann, Liadan Ita O'Callaghan, and Yoav Shoham. Truth revelation in approximately efficient combinatorial auctions. In *JACM 49(5)*, pages 577–602, Sept. 2002.

[29] Shahar Mendelson and Roman Vershynin. Entropy, combinatorial dimensions and random averages. In Jyrki Kivinen and Robert H. Sloan, editors, *COLT*, volume 2375 of *Lecture Notes in Computer Science*, pages 14–28. Springer, 2002.

[30] Noam Nisan and Amir Ronen. Computationally feasible vcg-based mechanisms. In *ACM Conference on Electronic Commerce*, 2000.

[31] Noam Nisan and Amir Ronen. Algorithmic mechanism design. *Games and Economic Behaviour*, 35:166 – 196, 2001. A preliminary version appeared in STOC 1999.

[32] Noam Nisan and Amir Ronen. Computationally feasible VCG-based mechanisms. *Journal of Artificial Intelligence Research*, 29:19–47, 2007. A preliminary version appeared in EC 2000.

[33] Noam Nisan and Ilya Segal. The communication requirements of e? cient allocations and supporting prices. *Journal of Economic Theory*, 129:192–224, 2006.

[34] Christos Papadimitriou, Michael Schapira, and Yaron Singer. On the hardness of being truthful. In *FOCS '08: Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science*, Philadelphia, PA, USA, 2008. IEEE Computer Society.

[35] Kevin Roberts. The characterization of implementable choice rules. In Jean-Jacques Laffont, editor, *Aggregation and Revelation of Preferences. Papers presented at the 1st European Summer Workshop of the Econometric Society*, pages 321–349. North-Holland, 1979.

[36] Tim Roughgarden. An algorithmic game theory primer. In *Proceedings of the 5th IFIP International Conference on Theoretical Computer Science (TCS). An invited survey.*, 2008.

[37] Norbert Sauer. On the density of families of sets. *J. Comb. Theory, Ser. A*, 13(1):145–147, 1972.

[38] Michael Schapira and Yaron Singer. Inapproximability of combinatorial public projects. In *WINE*, 2008.

[39] Saharon Shelah. A combinatorial problem; stability and order for models and theories in infinitary languages. *Pacific J Math*, 41:247–261, 1972.

[40] W. Vickrey. Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance*, pages 8–37, 1961.

[41] Jan Vondrák. Optimal approximation for the submodular welfare problem in the value oracle model. In Richard E. Ladner and Cynthia Dwork, editors, *STOC*, pages 67–74. ACM, 2008.

[42] Andrew Chi-Chih Yao. Some complexity questions related to distributive computing (preliminary report). In *STOC '79: Proceedings of the eleventh annual ACM symposium on Theory of computing*, pages 209–213, New York, NY, USA, 1979. ACM.

A Shattering Results

We first formally define the notion of "shattering" in a more general setting.

DEFINITION A.1. *For any sets U, V we interpret the notation V^U to mean the set of functions from U to V. If $R \subseteq V^U$, $S \subseteq U$, $L \subseteq V$, we say that S is (L,q)-shattered by R, for an integer q, $2 \leq q \leq |L|$, if there exist q functions $c_1, c_2, \ldots, c_q : S \to L$ that satisfy:*

1. $\forall x \in S \;\; \forall i \neq j \;\; c_i(x) \neq c_j(x)$
2. $\forall h \in [q]^S \;\; \exists f \in R \;\; \forall x \in S \;\; f(x) = c_{h(x)}(x)$

Intuitively, we associate with each element in S a range in L of size exactly q, and we say that S is (L,q)-shattered by R if every function that maps each element

in S to its associated range is a restriction of an element in R. In the context of combinatorial auctions, we see U as the set of items, and V as the set of bidders, plus a dummy bidder representing not allocating the item. Then set of functions V^U is the set of all possible allocations.

The following observation bridges this notion of shattering to its application to the combinatorial auctions in the paper.

OBSERVATION A.1. *If a subset S of size δm is (L,q)-shattered by $R \subseteq V^U$, then there exists a subset $L' \subseteq L$ and $S' \subseteq S$, such that $|L'| = q$, $|S'| \geq |S|/\binom{|L|}{q}$ and S' is (L',q)-shattered by R.*

The observation is easily seen by the pigeonhole principle. Note that by the definition of (L,q)-shattering, if $|L'| = q$, then we have that every function from S' to L' is a restriction of an element in R. In the context of combinatorial auctions, this means that all possible allocations of items in S' to the q bidders in L' are in the range R under restriction. It is this form of "strong" shattering that is in use in the main body of the paper. In the following lemmas, we will show the existence of large (L,q)-shattered sets, being aware that an application of the above observation implies a subset being "strongly" shattered, of size only a constant factor smaller.

LEMMA A.1. *For all integers $n \geq q \geq 2$, and every real number $\epsilon > 0$, there is a $\delta > 0$ such that the following holds. For every pair of finite sets M, N with $|N| = n$ and every set R of more than $(q-1+\epsilon)^{|M|}$ elements of N^M, there is a set S of at least $\delta|M|$ elements of M such that S is (V,q)-shattered by R.*

Proof. Let $F_q(m,n,d)$ denote the maximum cardinality of a set $R \subseteq A^B$ such that $|A| = n$, $|B| = m$, and R does not (A,q)-shatter any $(d+1)$-element subset of B.

Fix an element $b \in B$. For each element $f \in R$, let f_{-b} denote the restriction of f to the set $B\backslash\{b\}$. Take the set of all functions $g : B\backslash\{b\} \to A$ and partition it into sets $Q_0, Q_1, \cdots, Q_{\binom{n}{q}}$ as follows. First, given an ordered pair (g,a) consisting of a function g from $B\backslash\{b\}$ to A and an element $a \in A$, let $g * a$ denote the unique function f from B to A that maps b to a and restricts to g on $B\backslash\{b\}$. Now define $S(g)$ to be the set of all $a \in A$ such that $g * a$ is in R. Number all the q-element subsets of A from 1 to $\binom{n}{q}$, call them $P_1, P_2, \cdots, P_{\binom{n}{q}}$, and let Q_i $(1 \leq i \leq \binom{n}{q})$ consist of all g such that $S(g)$ has at least q elements, and the q smallest elements of $S(g)$ constitute P_i. Finally let Q_0 consist of all g such that $S(g)$ has fewer than q elements.

By our assumption that R does not (A,q)-shatter any set of size greater than d, we have the following facts:

1. Q_0 does not (A,q)-shatter any $(d+1)$-element subset of $B \setminus \{b\}$. Consequently,
$$|Q_0| \leq F(m-1,n,d).$$
2. For all $i \leq \binom{n}{q}$, Q_i does not (A,q)-shatter any d-element subset of $B \setminus \{b\}$. Consequently,
$$|Q_i| \leq F_q(m-1,n,d-1).$$

Let R_i denote the set of all $f \in R$ such that f_{-b} is in Q_i, for $0 \leq i \leq \binom{n}{q}$, then by definition of Q_i, we have $|R_0| \leq (q-1)|Q_0|$, and $|R_i| \leq n|Q_i|$ for $i \leq 1$. Since R_i's are disjoint, we have

$$|R| = \sum_{i=0}^{\binom{n}{q}} |R_i| \leq (q-1)|Q_0| + \sum_{i=1}^{\binom{n}{q}} n|Q_i|,$$

$$F_q(m,n,d) \leq \tag{A.1}$$

$$(q-1)F_q(m-1,n,d) + n\binom{n}{q}F_q(m-1,n,d-1)$$

The recurrence (A.1), together with the initial condition $F_q(m,n,0) = (q-1)^m$ for all m,n, implies the upper bound

$$F_q(m,n,d) \leq \sum_{i=0}^{d} n^i \binom{n}{q}^i \binom{m}{i}(q-1)^m$$

Thus, if $F_q(m,n,d) > (q-1+\epsilon)^m$ then, by using Stirling's approximation, we see that $d > \delta m$ for some δ depending only on ϵ and n.

In Section 4 of the paper, we made use of the fact that a range of allocations shatters a large subset if they generate good social welfare for many perfect valuations. The condition is captured by the following definition:

DEFINITION A.2. *For two functions $f, g \in N^M$, their* normalized Hamming distance $\mathrm{Ham}(f,g)$ *is equal to $\frac{1}{|M|}$ times the number of distinct $x \in M$ such that $f(x) \neq g(x)$. If $f \in N^M$ and $R \subseteq N^M$, the Hamming distance* $\mathrm{Ham}(f,R)$ *is the minimum of* $\mathrm{Ham}(f,g)$ *for all $g \in R$.*

As each perfect valuation can be seen as a function f in N^M, and each allocation can be viewed as a $g \in N^M$, $\mathrm{Ham}(f, g)$ is how much social welfare is lost by g on the perfect valuation f. In the same way, R can be viewed as a range of allocations, and $\mathrm{Ham}(f, R)$ is the minimum social welfare lost by any of the allocation in R on valuation f. If $\mathrm{Ham}(f, R)$ is small for a large fraction of $f \in N^M$, it means the range achieves a good approximation of social welfare for a significant portion of the perfect valuations.

We also note that since N can represent the set of bidders plus a dummy bidder representing not allocating an item, N^M can express all allocations including those not allocating all items. On the other hand, if we restrict the functions so that they can take values only in a subset L representing the real bidders, then they represent allocations that do not discard items. This explains the role played by the set L in the next lemma.

LEMMA A.2. *For every real number $\epsilon > 0$, every function $\gamma(n)$ bounded below by $1/\operatorname{poly}(n)$, and all integers $n \geq q \geq 2$, there is a $\delta > 0$ such that the following holds. For all finite sets M, N and all subsets $L \subseteq N$ with $|L| = n$, if $R \subseteq N^M$ and at least $\gamma(n)n^{|U|}$ points $f \in L^U$ satisfy $\mathrm{Ham}(f, R) < 1 - (q-1)/n - \epsilon$, then there is a set $S \subseteq M$ such that $|S| > \delta|M|$ and S is (L, q)-shattered by R.*

Proof. The proof parallels the counting argument in Section 3.1. Let $m = |M|, r = 1-(q-1)/n-\epsilon$. Let A be the set of all $f \in L^M$ such that $\mathrm{Ham}(f, R) < r$. Let G be a function from A to R such that $\mathrm{Ham}(f, G(f)) < r$ for all $f \in A$. Let $I(f)$ denote the set of all $x \in M$ such that $f(x) = G(f)(x)$. Our assumption that $\mathrm{Ham}(f, G(f)) < r$ implies that $|I(f)| \geq (\frac{q-1}{n} + \epsilon)m$. The number of pairs (f, J) such that $f \in A$, $|J| = \epsilon m/2$, $J \subseteq I(f)$ is bounded below by $\gamma(n)n^m \cdot \binom{(1/n+\epsilon)m}{\epsilon m/2}$. Henceforth we abbreviate $\gamma(n)$ as γ for convenience. By the pigeonhole principle, there is at least one set J of $\epsilon m/2$ elements such that the number of $f \in L^U$ satisfying $J \subseteq I(f)$ is at least

$$\begin{aligned}
&\gamma n^m \cdot \binom{(\frac{q-1}{n} + \epsilon)m}{\epsilon m/2} \Bigg/ \binom{m}{\epsilon m/2} \\
&= \gamma n^m \frac{((\frac{q-1}{n} + \epsilon)m)!\ ((1-\epsilon/2)m)!}{((\frac{q-1}{n} + \epsilon/2)m)!\ m!} \\
&> \gamma n^m \cdot \frac{(\frac{q-1}{n} + \epsilon)m}{m} \cdot \frac{(\frac{q-1}{n} + \epsilon)m - 1}{m-1} \cdots \frac{(\frac{q-1}{n} + \epsilon/2)m}{(1-\epsilon/2)m} \\
&> \gamma n^m \left(\frac{\frac{q-1}{n} + \epsilon/2}{1-\epsilon/2} \right)^{\epsilon m/2}.
\end{aligned}$$

Fix such a set J. For every $f \in L^M$ satisfying $J \subseteq I(f)$, the restriction of f to J is an element $g \in L^J$; note that g is also the restriction of $G(f)$ to J. For any single $g \in L^J$, the number of $f \in L^M$ that restrict to g is bounded above by $n^{m-\epsilon m/2}$. Applying the pigeonhole principle again, we see that the number of distinct $g \in L^J$ that occur as the restriction of some $f \in A$ satisfying $J \subseteq I(f)$ must be at least

$$\begin{aligned}
&\gamma n^m \left(\frac{\frac{q-1}{n} + \epsilon/2}{1-\epsilon/2} \right)^{\epsilon m/2} \Bigg/ n^{m-\epsilon m/2} \\
= \quad &\gamma \left(\frac{q-1+\epsilon n/2}{1-\epsilon/2} \right)^{\epsilon m/2}.
\end{aligned}$$

We now have the following situation. There is a set J of $\epsilon m/2$ elements, and at least $\gamma \cdot (q-1+\epsilon n/2)^{|J|}$ elements of L^J occur as the restriction of an element of R to J. It follows from Lemma A.1 that J has a subset of S of at least δm elements such that S is (L, q)-shattered by R.

Proof of Lemma 4.4: Combining Lemma A.1, Lemma A.2 and Observation A.1, we immediately get Lemma 4.4. □

B Omitted Proofs from Section 4.1

Proof of Lemma 4.1:

Suppose $S_1, \ldots, S_k$ are CD sets, with circuit families $\{\mathcal{C}_n^{(i)}\}$ $(1 \leq i \leq k)$ such that $\mathcal{C}_n^{(i)}$ has size bounded by a polynomial $q_i(n)$ and decides 3SAT correctly on all instances of size $n \in S_i$. Let $q(n)$ be a polynomial satisfying $q(n) \geq \max_{1 \leq i \leq k} q_i(n)$ for all $n \in \mathbb{N}$. We can obtain a family of circuits $\{\mathcal{C}_n\}$ of size bounded by $q(n)$, by defining $\mathcal{C}_n$ to be equal to $\mathcal{C}_n^{(i)}$ if n belongs to S_i but not to $S_1, \ldots, S_{i-1}$, and defining $\mathcal{C}_n$ to be arbitrary if $n \notin S_1 \cup \cdots \cup S_k$. Then $\mathcal{C}_n$ decides 3SAT correctly on all instances of size $n \in S_1 \cup \cdots \cup S_k$, as desired.

If $S_1, \ldots, S_k$ are CD sets, $p_1, \ldots, p_k$ are polynomials, and for $1 \leq i \leq k$ we have a PCD set $T_i \subseteq \bigcup_{n \in S_i} [n, p_i(n)]$, then we may take $p(n)$ to be any polynomial satisfying $p(n) \geq \max_{1 \leq i \leq k} p_i(n)$ for all $n \in \mathbb{N}$, and we may take S to be the set $S_1 \cup \cdots \cup S_k$. Then we find that the set $T = T_1 \cup \cdots \cup T_k$ is contained in $\bigcup_{n \in S} [n, p(n)]$. This implies that T is PCD, because S is CD. □

Proof of Lemma 4.2: By our assumption that $\mathcal{L}$ is NP-hard under polynomial-time many-one reductions, there is such a reduction from 3SAT to $\mathcal{L}$. Since the running time of the reduction is bounded by a polynomial $p(n)$, we know that it transforms a 3SAT instance of size n into an $\mathcal{L}$ instance of size at most $p(n)$. Assume without loss of generality that $p(n)$ is an increasing function of n.

Let S be the set of all n such that $\{p(n)+1, p(n)+2, \ldots, p(n+1)\}$ intersects T. The set S is complexity-defying, because for any $n \in S$ we can construct

a polynomial-sized circuit that correctly decides 3SAT instances of size n, as follows. First, we take the given 3SAT instance and apply the reduction from the preceding paragraph to transform it into an $\mathcal{L}$ instance of size at most $p(n)$. Then, letting m be any element of $T \cap \{p(n)+1, \ldots, p(n+1)\}$, we apply the padding reduction to transform this $\mathcal{L}$ instance into another $\mathcal{L}$ instance of size m. Finally, we solve this instance using a circuit of size poly(m) that correctly decides $\mathcal{L}$ on all instances of size m; such a circuit exists by our assumption on T.

For every $m \in T$ there is an $n \in \mathbb{N}$ such that $p(n) < m \leq p(n+1)$, and this n belongs to S. Thus, $T \subseteq \bigcup_{n \in S} [n, p(n+1)]$, and this confirms that T is PCD. □

Proof of Lemma 4.3: Suppose that

$$\mathbb{N} \subseteq \bigcup_{n \in S} [n, p(n)] \tag{B.2}$$

for some complexity-defying set S and polynomial function $p(n)$. We may assume without loss of generality that $p(n)$ is an increasing function of n and that $p(n) \geq n$ for all n.

Suppose that $\{\mathcal{C}_n\}$ is a polynomial-sized circuit family that correctly decides 3SAT whenever the input size is in S. We will construct a polynomial-sized circuit family that correctly decides 3SAT on all inputs. The construction is as follows: given an input size m, using (B.2) we may find a natural number n such that $n \leq p(m) \leq p(n)$. Since p is an increasing function, we know that $n \geq m$. Given an instance of 3SAT of size m, we first adjoin irrelevant clauses that don't affect its satisfiability — e.g. the clause $(x \vee \overline{x})$ — until the input size is increased to n. This transformation can be done by a circuit of size poly(m), since $n \leq p(m)$. Then we solve the new 3SAT instance using the circuit $\mathcal{C}_n$. By our assumption on S, this correctly decides the original 3SAT instance of size m. As m was arbitrary, this establishes that $\mathsf{NP} \subseteq \mathsf{P}/\,\mathrm{poly}$, as desired. □